arXiv:1511.06918v1 [cs.GT] 21 Nov 2015

Ironing in the Dark

Tim Roughgarden* Okke Schrijvers†

November 24, 2015

Abstract

This paper presents the first polynomial-time algorithm for position and matroid auction environments that learns, from samples from an unknown bounded valuation distribution, an auction with expected revenue arbitrarily close to the maximum possible. In contrast to most previous work, our results apply to arbitrary (not necessarily regular) distributions and the strongest possible benchmark, the Myerson-optimal auction. Learning a near-optimal auction for an irregular distribution is technically challenging because it requires learning the appropriate "ironed intervals," a delicate global property of the distribution.

1 Introduction

The traditional economic approach to revenue-maximizing auction design, exemplified by Myerson [27], posits a known prior distribution over what bidders are willing to pay, and then solves for the auction that maximizes the seller's expected revenue with respect to this distribution. Recently, there has been an explosion of work in computer science that strives to make the classical theory more "data-driven," replacing the assumption of a known prior distribution with that of access to relevant data, in the form of samples from an unknown distribution. In this paper, we study the problem of learning a near-optimal auction from samples, adopting the formalism of Cole and Roughgarden [9]. The idea of the model, inspired by PAC-learning [32], is to parameterize through samples the "amount of knowledge" that the seller has about bidders' valuation distributions.

We consider single-parameter settings, where each of n bidders has a private valuation (i.e., willingness to pay) for "winning" and valuation 0 for "losing." Feasible outcomes correspond to subsets of bidders that can simultaneously win; the feasible subsets are known in advance.[1] We assume that bidders' valuations are drawn i.i.d. from a distribution F that is unknown to the seller. However, we assume that the seller has access to m i.i.d. samples from the distribution F — for example, bids that were observed in comparable auctions in the past. The goal is to design a polynomial-time algorithm $A(v_1, \ldots, v_m)$ from samples $v_i \sim F$ to truthful auctions that, for every

*Department of Computer Science, Stanford University, 474 Gates Building, 353 Serra Mall, Stanford, CA 94305. This research was supported in part by NSF grant CCF-1215965. Email: tim@cs.stanford.edu.

†Department of Computer Science, Stanford University, 482 Gates Building, 353 Serra Mall, Stanford, CA 94305. This research was supported in part by NSF grant CCF-1215965. Email: okkes@cs.stanford.edu.

[1] For example, in auction with k copies of an item, where each bidder only wants one copy, feasible outcomes correspond to subsets of at most k bidders.

distribution F, achieves expected revenue at least $1-\epsilon$ times the optimal expected revenue.[2] The sample complexity of achieving a given approximation factor $1-\epsilon$ is then the minimum number of samples m such that there exists a learning algorithm A with the desired approximation. This model serves as potential "sweet spot" between worst-case and average-case analysis, inheriting much of the robustness of the worst-case model (since we demand guarantees for every underlying distribution F) while allowing very good approximation guarantees.

1.1 Our Results

We give polynomial-time algorithms for learning a $(1-\epsilon)$-approximate auction from samples, for arbitrary matroid (or position) auction environments and arbitrary (not necessary regular) distributions with bounded valuations. (See Section 2 for definitions.)

Precisely, our main result is a polynomial-time algorithm that, given a matroid environment with n bidders and m i.i.d. samples from an arbitrary distribution F with maximum valuation H, with probability $1-\delta$, approximates the maximum-possible expected revenue up to an additive loss of at most $3n\sqrt{\frac{\ln 2\delta^{-1}}{2m}} \cdot H$. Thus for every $\epsilon > 0$, the additive loss is at most ϵ (with probability at least $1-\delta$) provided $m = \Omega(n^2 H^2 \epsilon^{-2} \log \delta^{-1})$. When valuations lie in $[1, H]$, so that the optimal expected revenue is bounded below, this result immediately implies the same sample complexity bound for learning a $(1-\epsilon)$-(multiplicative) approximate auction. Our main result can also be used to give a no-regret guarantee in a stochastic learning setting (Section 5).

A lower bound of Cesa-Bianchi et al. [6] implies that, already for simpler settings, the quadratic dependence of our sample complexity bound on $\frac{1}{\epsilon}$ is optimal. A lower bound of Huang et al. [22] implies that, already with a single bidder, the sample complexity must depend polynomially on H. Whether or not the sample complexity needs to depend on n is an interesting open question.

Our approach is based on a "swiching trick" (Proposition 2.1) and we believe it will lead to further applications. A key idea is to express the difference in expected revenue between an optimal and a learned auction in a matroid or position environment purely in terms of a difference in area under the true and estimated revenue curves. This "global" analysis avoids having to compare explictly the true and estimated virtual valuations or the optimal and learned allocation rules. With this approach, there is clear motivation behind each of the steps of the learning algorithm, and the error analysis, while non-trivial, remains tractable in settings more general than those studied in most previous works.

1.2 Why Irregular Distributions Are Interesting

A major difference between this paper and much previous work on learning near-optimal auctions from samples is that our positive results accommodate *irregular distributions*. To explain, we first note that learning a near-optimal auction from samples is impossible without some restriction on the space of possible valuation distributions F.[3] This observation motivates imposing the weakest-possible conditions under which non-trivial positive results are possible.

[2] By a truthful auction, we mean one in which truthful bidding is a dominant strategy for every bidder. The restriction to dominant strategies is natural given our assumption of an unknown distribution. Given this, the restriction to truthful auctions is without loss of generality (by the "Revelation Principle," see [29]). Also, for the single-parameter problems that we study, there is a truthful optimal auction.

[3] To see this, consider a single-bidder problem and all distributions that take on a value M^2 with probability $\frac{1}{M}$ and 0 with probability $1-\frac{1}{M}$. The optimal auction for such a distribution earns expected revenue at least M. It is not difficult to prove that, for every m, there is no way to use m samples to achieve near-optimal revenue for

A majority of the literature on approximation guarantees for revenue maximization (via learning algorithms or otherwise) restricts attention to "regular" valuation distributions or subclasses thereof; see related work below for examples and exceptions. Formally, a distribution F with density f is *regular* if

$$\varphi(v) = v - \frac{1 - F(v)}{f(v)} \tag{1}$$

is a nondecreasing function of v. φ is also called the *virtual valuation* function. Intuitively, regularity is a logconcavity-type assumption that provides control over the tail of the distribution. While many important distributions are regular, plenty of natural distributions are not. For example, Sivan and Syrgkanis [31] point out that mixtures of distributions (even of uniform distributions) tend to be irregular, and yet obviously occur in the real world. This motivates relaxing regularity conditions while still somehow excluding pathological distributions. One interesting approach is to parameterize a distribution by its "degree of irregularity;" see [20, 22, 31] for some proposals of how to do this. Another approach, and the one we take here, is to bound the support of a distribution (to $[0, H]$ or $[1, H]$, say) and allow it to be otherwise arbitrary. Bounded but otherwise arbitrary valuation distributions clearly capture interesting settings beyond those modeled by regular distributions.

1.3 Why Irregular Distributions Are Hard

To understand why irregular distributions (with bounded valuations) are so much more technically challenging than regular distributions, we need to review some classical optimal auction theory. We can illustrate the important points already in single-item auctions. Myerson [27] proved that, for every regular distribution F, the optimal auction is simply a second-price auction supplemented with a reserve price of $\varphi^{-1}(0)$, where φ denotes the virtual valuation function in (1). (The winner, if any, pays the higher of the reserve price and the second-highest bid.) Thus, *learning the optimal auction reduces to learning the optimal reserve price*, a single statistic of the unknown distribution. And indeed, for an unknown regular distribution F, there is a polynomial-time learning algorithm that needs only $\text{poly}(\frac{1}{\epsilon})$ samples to compute a $(1-\epsilon)$-approximate auction [12, 22].

The technical challenge of irregular distributions is the need to *iron*. When the virtual valuation function φ of the distribution F is not nondecreasing, Myerson [27] gave a recipe for transforming φ into a nondecreasing "ironed" virtual valuation function $\overline{\varphi}$ such that the optimal single-item auction awards the item to the bidder with the highest positive ironed virtual valuation (if any), breaking ties randomly. Intuitively, this ironing procedure identifies intervals of non-monotonicity in φ and changes the value of the function to be constant on each of these intervals. (See also below and the exposition by Hartline [20].)

The point is that the appropriate ironing intervals of a distribution are a *global* property of the distribution and its (unironed) virtual valuation function. Estimating the virtual valuation function at a single point — all that is needed in the regular case — would appear much easier than estimating the right intervals to iron in the irregular case.

We present two examples to drive this point home. The first, which is standard, shows that not ironing can lead to a constant-factor loss in expected revenue. The second shows that tiny mistakes

every such distribution — for sufficiently large M, all m samples are 0 w.h.p. and the algorithm has to resort to an uneducated guess for M.

in the choice of ironing intervals can lead to a large loss of expected revenue.[4]

Example 1.1 (Ironing Is Necessary for Near-Optimal Revenue) The distribution is as follows: with probability $1/H$ the value is H (for a large H) and it is 1 otherwise. The optimal auction irons the interval $[1, H)$ for expected revenue of $2-\frac{1}{n}$ [20], which approaches 2 with many bidders n. Auctions that do not implicitly or explicitly iron obtain expected revenue only 1.

Example 1.2 (Small Mistakes Matter) Let F be 5 with probability 1/10 and 1 otherwise, and consider a single-item auction with 10 bidders. The optimal auction irons the interval $[1, 5)$ and has no reserve price. If there are at least two bidders with value 5 one of them will get the item at price 5; if all bidders have value 1, one of them will receive it at price 1. If there is exactly one bidder with value 5, then her price is $\frac{1}{10} \cdot 1 + \frac{9}{10} \cdot 5 = \frac{46}{10}$.

Now consider an algorithm that slightly *overestimated* the end of the ironing interval to be $[1, 5 + \epsilon)$ with $\epsilon > 0$. (Imagine F actually has small but non-zero density above 5, so that this mistake could conceivably occur.) Now all bids always fall in the ironing interval and therefore the item is always awarded to one of the players at price 1. Not only do we lose revenue when there is exactly one bidder, but additionally we lose revenue for auctions with at least two bidders with value 5. This auction has even worse revenue than the standard second-price auction, so the attempt to iron did more harm than good.

If we adapt Example 1.2 to slightly underestimate the ironing interval, there is almost no loss in revenue. This motivates an important idea in our learning algorithm, which is to closely protect against overestimation.

1.4 Related Work

Elkind [16] gives a polynomial-time learning algorithm for the restricted case of single-item auctions with discrete distributions with known finite supports but with unknown probabilities. In the model in [16], learning is done using an oracle that compares the expected revenue of pairs of auctions, and $O(n^2K^2)$ oracles calls suffice to determine the optimal auction (where n is the number of bidders and K is the support size of the distributions). Elkind [16] notes that such oracle calls can be implemented approximately by sampling (with high probability), but no specific sample complexity bounds are stated.

Cole and Roughgarden [9] also give a polynomial-time algorithm for learning a $(1-\epsilon)$-approximate auction for single-item auctions with possibly non-identical bidders, under incomparable assumptions to [16]: valuation distributions that can be unbounded but must be strongly regular. It is necessary and sufficient to have $m = \text{poly}(n, \frac{1}{\epsilon})$ samples, however in the analysis in [9] the exponent in the upper bound is large (currently, 10).

The papers of Cesa-Bianchi et al. [5] and Medina and Mohri [25] give algorithms for learning the optimal reserve-price-based single-item auction. Recall from Example 1.1 that, with irregular distributions, the best reserve-price-based auction can have expected revenue far from optimal.

Dughmi et al. [13] proved negative results (exponential sample complexity) for learning near-optimal mechanisms in multi-parameter settings that are much more complex than the single-parameter settings studied here. The paper also contains positive results for restricted classes of mechanisms.

[4] See Appendix A for a different example that demonstrates why irregular distributions are harder than regular ones.

Huang et al. [22] give optimal sample complexity bounds for the special case of a single bidder under several different distributional assumptions, including for the case of bounded irregular distributions where they need $O(H \cdot \epsilon^{-2} \cdot \log(H\epsilon^{-1}))$ samples.

Morgenstern and Roughgarden [26] recently gave general sample complexity upper bounds which are similar to ours and cover all single-parameter settings (matroid and otherwise), although the (brute-force) learning algorithms in [26] are not computationally efficient.

For previously studied models about revenue-maximization with an unknown distribution, which differ in various respects from the model of Cole and Roughgarden [9], see [2, 6, 23]. For other ways to parameterize partial knowledge about valuations, see e.g. [1, 8]. For other uses of samples in auction design that differ from ours, see Fu et al. [17], who use samples to extend the Crémer-McLean theorem [10] to partially known valuation distributions, and Chawla et al. [7], which is discussed further below. For asymptotic optimality results in various symmetric settings (single-item auctions, digital goods), which identify conditions under which the expected revenue of some auction of interest (e.g., second-price) approaches the optimal with an increasing number of i.i.d. bidders, see Neeman [28], Segal [30], Baliga and Vohra [4], and Goldberg et al. [18]. For applications of learning theory concepts to prior-free auction design in unlimited-supply settings, see Balcan et al. [3].

Finally, the technical issue of ironing from samples comes up also in Ha and Hartline [19] and Chawla et al. [7], in models incomparable to the one studied here. The setting of [19] is constant-factor approximation guarantees for prior-free revenue maximization, where the goal is input-by-input rather than distribution-by-distribution guarantees. Chawla et al. [7] study non-truthful auctions, where bidders' true valuations need to be inferred from equilibrium bids, and aim to learn the optimal "rank-based auction," which can have expected revenue a constant factor less than that of an optimal auction. Our goal of obtaining a $(1-\epsilon)$-approximation of the maximum revenue achieved by any auction is impossible in the settings of [19, 7].

Summarizing, this paper gives the first polynomial-time algorithm for position and matroid environments that learns, from samples from an unknown irregular valuation distribution, an auction with expected revenue arbitrarily close to the maximum possible.

1.5 Organization

Section 2 covers learning and auction preliminaries important for our analysis. Section 3 presents many of our key ideas in the special case of a single-item auction. Section 4 extends our results to position and matroid auction environments. Section 5 deduces a no-regret guarantee from our learning algorithms.

2 Preliminaries

2.1 Empirical Cumulative Distribution Function and the DKW Inequality

Let $X = \{X_i\}_{i=1}^m$ be a set of m samples, and let $X^{(i)}$ be the i^{th} order statistic. We use the standard notion of the empirical cumulative distribution function (empirical CDF): $\widehat{F_m}(v) = \frac{1}{m} \cdot |\{X_i : X_i \leq v\}|$. The empirical CDF is an estimator for the quantile of a given value. The Dvoretzky-Kiefer-Wolfowitz (DKW) inequality [14, 24] states that the difference between the empirical CDF and the actual CDF decreases quickly in the number of samples. Let $\epsilon_{m,\delta} = \sqrt{\frac{\ln 2\delta^{-1}}{2m}}$, then

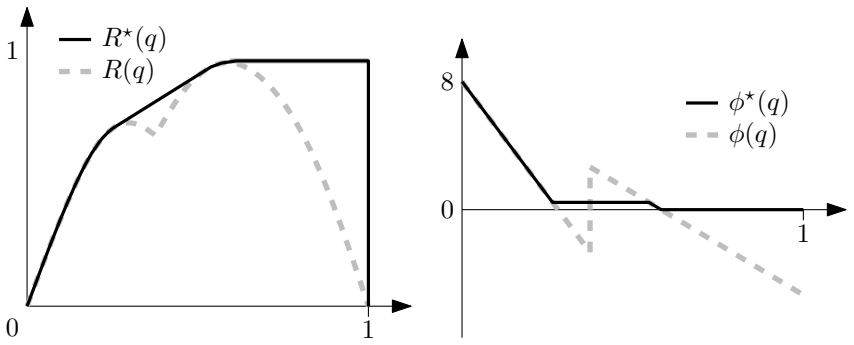

(a) Revenue curve $R(q)$ and the optimal induced revenue curve $R^\star(q)$.

(b) Virtual value function $\varphi(q)$ and the induced virtual value function $\varphi^\star(q)$.

Figure 1: The dashed gray line is the revenue curve $R(q)$ and its derivative, the virtual value function $\varphi(q)$, for a irregular bimodal distribution. In black we have the curves φ^{opt} and R^{opt} corresponding to the optimal ironing interval, and optimal reserve price for this distribution. This example taken from [20, Chapter 3].

$\Pr\left[\sup_{v\in\mathbb{R}}\left|F(v)-\widehat{F_m}(v)\right|\le\epsilon_{m,\delta}\right]\ge 1-\delta$. So the largest error in the empirical CDF shrinks as $O(m^{-1/2})$. For our purposes we will not need the CDF F, but rather its inverse F^{-1}. Define $\widehat{F_m^{-1}}(x)$ as $X^{(\max(1,\lceil x\cdot m\rceil))}$ for $x\in[0,1]$. (For convenience, define $\widehat{F_m^{-1}}(x)$ as 0 if $x<0$ and H if $x>1$.) By definition, for all $v\in[0,H]$:

$$\widehat{F_m^{-1}}\left(\widehat{F_m}(v)\right) \quad\le\quad v \quad\le\quad \widehat{F_m^{-1}}\left(\widehat{F_m}(v)+\frac{1}{m}\right). \tag{2}$$

In the remainder of this paper, we will use $\widehat{F}$, $\widehat{F^{-1}}$, and ϵ without explicitly referring to the number of samples m and confidence parameter δ.

2.2 Optimal Auctions using the Revenue Curve

Revenue Curve and Quantile Space For a value v and a distribution F, we define the corresponding *quantile* as $1-F(v)$.[5] The quantile of v corresponds to the sale probability for a single bidder with valuation drawn from F as the posted price v. Note that higher values correspond to lower quantiles. The *revenue curve* R for value distribution F is the curve, in quantile space, that is defined as $R(q)=q\cdot F^{-1}(1-q)$, see the dashed curve in Figure 1a. Myerson [27] showed that the ex ante expected revenue for a bidder is $\mathbb{E}_{q\sim U[0,1]}[-R(q)\cdot y'(q)]$, where y is the interim allocation function for the standard Vickrey auction (defined formally below). The derivative of R is the virtual value function $\varphi(.)$ — recall (1) — in quantile space, see Figure 1b. For regular distributions the revenue curve is concave and the virtual value function non-increasing, but for irregular distributions this is not the case.

Myerson [27] showed that for any non-concave revenue curve R, one can create an allocation rule that will yield the same revenue as R's convex hull $\mathcal{CH}(R)$. This procedure is called ironing, and for each interval where R differs from $\mathcal{CH}(R)$, we take the virtual value to be slope of the convex hull over this interval. This means the virtual values are constant, and hence any two bids

[5] This is for consistency with recent literature; "reversed quantile" might be a more accurate term.

in that range are interchangeable, so the ironed allocation function is constant on this interval.[6] The resulting *ironed* revenue curve $R^\star$ will be concave and the corresponding ironed virtual value function $\varphi^\star$ will be monotonically non-increasing. It is also useful to think of a reserve price r (with corresponding quantile q_r) in a similar way, as effectively changing the virtual valuation function so that $\varphi^\star(q) = 0$ whenever $q \geq q_r$ (Figure 1b), with the corresponding revenue curve $R^\star$ constant in that region (Figure 1a).[7]

More generally, given any set of disjoint ironing intervals $\mathcal{I}$ and reserve price r, both in value space, we can imagine these modifying the revenue curve as follows. (For now this is a thought experiment; Proposition 2.1 connects these revenue curve modifications back to allocation rule modifications.) Let R be the revenue curve without ironing or a reserve price, and define $R^{(\mathcal{I},r)}$ as the revenue curve *induced* by a set $\mathcal{I}$ of ironing intervals and reserve price r. This curve is defined by

$$R^{(\mathcal{I},r)}(q) = \begin{cases} R(F(r)) & \text{if } q > F(r) \\ R(q_a) + \frac{q-q_a}{q_b-q_a}\left(R(q_b) - R(q_a)\right) & \text{if } q \in (q_a, q_b] \text{ with } [F^{-1}(1-q_b), F^{-1}(1-q_a)) \in \mathcal{I} \\ R(q) & \text{otherwise.} \end{cases} \tag{3}$$

Given $\mathcal{I}$ and r as above, we define the auction $A^{(\mathcal{I},r)}$ as follows: given a bid profile: (i) reject every bid with $b_i < r$; (ii) for each ironing interval $[a, b) \in \mathcal{I}$, treat all bids $\{b_j : a \leq b_j < b\}$ as identical (equal to some common number between a and b); (iii) among the remaining bidders, maximize the social welfare (sum of the ironed bids of the winners); (iv) charge payments so that losers always pay 0 and so that truthful bidding is a dominant strategy for every player. This auction is well defined (i.e., independent of the choice of the common numbers in (ii)) in settings where the computation in (iii) depends only on the ordering of the ironed bids, and not on their numerical values. In this case, the payments in (iv) are uniquely defined (by standard mechanism design results). This is the case in every matroid environment[8] and also in position auctions (see [15, 33]). In such a setting, we use $\mathcal{A}$ to denote the set of all auctions of the form $A^{(\mathcal{I},r)}$. We restrict attention to such settings for the rest of the paper.

The Switching Trick Given a distribution F, we explained two ways to use ironing intervals $\mathcal{I}$ and a reserve price r: (i) to define a modified revenue curve $R^{(\mathcal{I},r)}$ (and hence virtual valuations); or (ii) to define an auction $A^{(\mathcal{I},r)}$. The "switching trick" formalizes the connection between them: the expected virtual welfare of the welfare-maximizing auction with the modified virtual valuations (corresponding to the derivative of $R^{(\mathcal{I},r)}$) equals the expected virtual welfare of the modified auction $A^{(\mathcal{I},r)}$ with the original virtual valuations.

More formally, let $x_i : \mathbb{R}_+^n \to \mathbb{R}_+$ be the ex-post allocation function of the welfare maximizing truthful auction that takes the bids $\mathbf{b}$ of all players and results in the allocation to bidder i. The

[6] It takes some effort to show that keeping the allocation probability constant on an interval has exactly the effect we described here [27].

[7] Most of the existing literature would not consider the effect of the reserve price on the revenue curve, in which case the black and dashed lines would coincide after the second peak. However, by including its effect as we did, we'll be able to apply the Switching Trick described below.

[8] In a matroid environment, the set $\mathcal{F}$ of feasible outcomes satisfies: (i) (downward-closed) $T \in \mathcal{F}$ and $S \subseteq T$ implies $S \in F$; and (ii) (exchange property) whenever $S, T \in \mathcal{I}$ with $|T| < |S|$, there is some $i \in S \setminus T$ such that $T \cup \{i\} \in \mathcal{I}$.

interim allocation function $y_i : [0,1] \to \mathbb{R}_+$ is the expected allocation to bidder i when her quantile is q, where the expectation is over the quantiles of the other bidders: $y_i(q_i) = \mathbb{E}_{\mathbf{q}_{-i} \sim U[0,1]^{n-1}}[x(F^{-1}(1-q), F^{-1}(\mathbf{1} - \mathbf{q}_{-i}))]$ where $F^{-1}(\mathbf{1} - \mathbf{q}_{-i})$ is $\mathbf{b}_{-i}$ for which each $b_j = F(1-q_j)$. For example, in the standard Vickrey (single-item) auction with n bidders, every bidder i has the interim allocation function $y_i(q) = (1-q)^{n-1}$.[9]

For every auction of the form $A^{(\mathcal{I},r)}$, the interim allocation function $y_i^{(\mathcal{I},r)}$ of a bidder i can be expressed in terms of the interim allocation function y_i without ironing and reserve price (see also Figure 1b):

$$y_i^{(\mathcal{I},r)}(q) = \begin{cases} 0 & \text{if } q > F(r) \\ \frac{1}{q_b - q_a} \int_{q_a}^{q_b} y(q)dq & \text{if } q \in [q_a, q_b) \text{ with } [F^{-1}(1-q_b), F^{-1}(1-q_a)) \in \mathcal{I} \\ y_i(q) & \text{otherwise.} \end{cases} \tag{4}$$

Proposition 2.1 (Switching Trick) *Consider a matroid or position auction setting, as above. For every valuation distribution F, every reserve price r, every set $\mathcal{I}$ of disjoint ironing intervals, and every bidder i,*

$$\mathbb{E}_{q \sim U[0,1]}[R(q) \cdot (y_i^{(\mathcal{I},r)})'(q)] = \mathbb{E}_{q \sim U[0,1]}[R^{(\mathcal{I},r)}(q) \cdot y_i'(q)].$$

Proof: Fix F, $\mathcal{I}$, r, and y. Let $y^{(\mathcal{I},r)}$ be the interim allocation rule from running auction $A_{\mathcal{I},r}$. Let R be the revenue curve of F and let $R^{(\mathcal{I},r)}$ denote the revenue curve induced by $\mathcal{I}$ and r.

- Define a distribution $F^{(\mathcal{I},r)}$ (which is not equal to F unless $\mathcal{I} = \emptyset$ and $r = 0$) that has the property that its revenue curve $q \cdot F^{(\mathcal{I},r)}(1-q)$ is $R^{(\mathcal{I},r)}$. To see that this is well-defined, observe the following. Any line ℓ through the origin only intersects $R^{(\mathcal{I},r)}$ once (if there are point masses in F then a line through the origin intersect R in a single interval). This means that we can use $R^{(\mathcal{I},r)}$ to construct $F^{(\mathcal{I},r)}$: $F^{(\mathcal{I},r)}(v)$ is the q for which $q \cdot v$ intersects with $R^{(\mathcal{I},r)}(q)$ (if there are any point masses then there will be a range of q for which this is the case; in that case take the largest such q). Alternatively, see Hartline and Roughgarden [21] for an explicit formula for $F^{(\mathcal{I},r)}$.

- If we run the same auction $A_{\mathcal{I},r}$ on bidders with values drawn from $F^{(\mathcal{I},r)}$, the expected revenue is identical to the auction with bidder values drawn from F:

  $$\mathbb{E}_{q \sim U[0,1]}[R(q) \cdot \left(y^{(\mathcal{I},r)}\right)'(q)] = \mathbb{E}_{q \sim U[0,1]}[R^{(\mathcal{I},r)}(q) \cdot \left(y^{(\mathcal{I},r)}\right)'(q)].$$

  This can easily be seen by filling in the definitions from (3) and (4).

- If the bidders have distribution $F^{(\mathcal{I},r)}$, then we might as well not iron or have a reserve price at all; so

  $$\mathbb{E}_{q \sim U[0,1]}[R^{(\mathcal{I},r)}(q) \cdot \left(y^{(\mathcal{I},r)}\right)'(q)] = \mathbb{E}_{q \sim U[0,1]}[R^{(\mathcal{I},r)}(q) \cdot y'(q)].$$

  This is also easily seen by filling in the definitions.

■

[9] In general matroid settings, different bidders can have different interim allocation functions (even though valuations are i.i.d.)

Revenue Curve	Description
R	$q \cdot F^{-1}(1-q)$
$\widehat{R}_{\min}$	$q \cdot \widehat{F^{-1}}(1-q-\epsilon)$
$\widehat{R}_{\max}$	$q \cdot \widehat{F^{-1}}(1-q+\epsilon+\frac{1}{m})$

Table 1: Overview of notation for revenue curves.

2.3 Notation

In the remainder of this paper, our analysis will rely on bounding the difference in revenue of an auction with respect to the optimal auction in terms of their revenue curves. We will use the following conventions, see Table 1. The unaltered revenue curve for distribution is denoted by $R(q) = q \cdot F^{-1}(1-q)$. To denote when we use an estimator for a revenue, i.e. a revenue curve that is constructed based on samples, we use a hat: $\widehat{R}(q) = q \cdot \widehat{F^{-1}}(1-q)$. Based on the available samples we construct high-probability upper and lower bounds for R, that are thus denoted as $\widehat{R}_{\max}(q) = q \cdot \widehat{F^{-1}}(1-q+\epsilon+\frac{1}{m})$ and $\widehat{R}_{\min}(q) = q \cdot \widehat{F^{-1}}(1-q-\epsilon)$.

We use superscript to denote when a revenue curve is ironed and has a reserve price. For a general set of ironing intervals $\mathcal{I}$ and reserve price r, $R^{(\mathcal{I},r)}$ is the revenue curve induced by it, see (3). The superscript $\star$ denotes that the revenue curve is optimally ironed and reserved, i.e. $R^\star$ is the revenue curve of Myerson's auction using F, and $\widehat{R}^\star_{\max}$ is the revenue curve corresponding to the convex hull of $\widehat{R}_{\max}$ that additionally stays constant after the highest point. Finally, we'll use R^{alg} and R^{opt} to denote an algorithm ALG's revenue curve and the optimal revenue curve for F respectively (thus $R^{opt} = R^\star$, but we'll use R^{opt} to emphasize its relation to R^{alg}).

For the ironing intervals $\mathcal{I}$ (and reserve price r) we use $\mathcal{I}_q$ (resp. r_q) when it is important that the ironing intervals are defined in *quantile space*. Finally, $\mathcal{I}_{opt}$, $\mathcal{I}_{alg}$, and $\mathcal{I}_{max}$ (and similarly for reserve r) refer to the ironing intervals of the optimal auction, algorithm ALG and the optimal ironing intervals for $\widehat{R}_{\max}$ respectively.

3 Additive Loss in Revenue for Single-Item Auctions

In this section we describe an algorithm that takes a set X of m samples, and a confidence parameter δ as input, and outputs a set $\mathcal{I}$ of ironing intervals and a reserve price r, both in value space. We focus on the case where $\mathcal{I}$ and r are used in a single-item auction $A^{\mathcal{I},r}_{(1)} \in \mathcal{A}$ (recall the notation in Section 2) and show that the additive loss in revenue of $A^{\mathcal{I},r}_{(1)}$ with respect to the revenue of the optimal auction $A^{opt}_{(1)}$ for single-item auctions is $O(\epsilon \cdot n \cdot H)$, with $\epsilon = \sqrt{\frac{\ln 2\delta^{-1}}{2m}}$. In section 4 we extend the results to matroid and position auctions.

Theorem 3.1 (Main Theorem) *For a single-parameter environment with optimal auction of the form $A^{(\mathcal{I},r)}$ with n i.i.d. bidders with values from unknown irregular distribution F, m i.i.d. samples from F, with probability $1-\delta$, the additive loss in expected revenue of Algorithm 2 compared to the optimal expected revenue is at most $3 \cdot n \cdot H \cdot \sqrt{\frac{\ln 2\delta^{-1}}{2m}}$.*

Algorithm 1 Compute the ironing intervals $\mathcal{I}$ and reserve price r.

ComputeAuction(X, δ)

1 Construct $\widehat{F^{-1}}$ from X; let $\epsilon = \sqrt{\frac{\ln 2 \cdot \delta^{-1}}{2|X|}}$.
2 Construct $\widehat{R}_{\min}(q) = q \cdot \widehat{F^{-1}}(1 - q - \epsilon)$.
3 Compute the convex hull $\mathcal{CH}(\widehat{R}_{\min})$, of $\widehat{R}_{\min}$.
4 Let $\mathcal{I}_q$ be the set of intervals where $\widehat{R}_{\min}$ and $\mathcal{CH}(\widehat{R}_{\min})$ differ.
5 **for each** quantile ironing interval $(a_i, b_i) \in \mathcal{I}_q$
6 Add $[\widehat{F^{-1}}(1 - b_i - \epsilon),\ \widehat{F^{-1}}(1 - a_i - \epsilon))$ to $\mathcal{I}$.
7 Let the reserve quantile be $r_q = \arg\max_q \widehat{R}_{\min}(q)$.
8 Let the reserve price be $r = \widehat{F^{-1}}(1 - r_q - \epsilon)$.
9 **return** $(\mathcal{I}, r)$

Algorithm 2 Empirical Myerson.

EmpiricalMyerson$(X, \delta, \mathbf{b})$

1 $\mathcal{I}, r \leftarrow$ ComputeAuction(X, δ)
2 **return** $A_{(1)}^{\mathcal{I},r}(\mathbf{b})$

3.1 The Empirical Myerson Auction

We run the Empirical Myerson auction (Algorithm 2), which we divided this into two parts: the first is an algorithm *ALG* (Algorithm 1) that computes the ironing intervals $\mathcal{I}$ and reserve price r based on samples X and confidence parameter δ. The second step is to run the welfare-maximizing auction subject to ironing and reservation. In this section we focus on analyzing the single-item auction, but the only place this is used is in line 2 of Algorithm 2. Auctions for position auctions and matroid environments are defined completely analogously.

The Empirical Myerson auction takes an estimator for the quantile function $\widehat{F^{-1}}$ and constructs its revenue curve.[10] From this, the convex hull $\mathcal{CH}(R)$ is computed and wherever $\mathcal{CH}(R)$ and R disagree, an ironing interval is placed. Then, the highest point on R is used to obtain the reserve price quantile $q_r = \arg\max_q R(q)$. Note that this is all done in quantile space, but we need to specify the reserve price and ironing intervals in value space. So the last step is to use the empirical CDF $\widehat{F}$ to obtain the values at which to place the reserve price and ironing intervals.

Algorithm 2 follows that approach, with the exception that in line 2 of ComputeAuction, we take the empirical quantile function to be $\widehat{F^{-1}}(1 - q - \epsilon)$ rather than the arguably more natural choice of $\widehat{F^{-1}}(1 - q)$. By the DKW inequality, $\widehat{F^{-1}}(1 - q - \epsilon) \leq F^{-1}(1 - q)$ with probability $1 - \delta$ (we prove this in Lemma 3.3). The motivation here is to protect against overestimation — recall

[10] There are variations that differ from this approach, most notably the Random Sampling Empirical Myerson (RSEM) [11] which constructs the revenue curve from the bidders in the auction, rather than from past data. This leads to constant-factor approximation bounds for prior-free auctions.

the cautionary tale of Example 1.2. That this approach indeed leads to good revenue guarantees is shown in this section.

3.2 Additive Revenue Loss in Terms of Revenue Curves

We start with a technical lemma that reduces bounding the loss in revenue to bounding the estimation error due to using samples as opposed to the true distribution F.

Lemma 3.2 *For a distribution F, let R^{alg} be the revenue curve induced by an algorithm $ALG \in \mathcal{A}$ and let R^{opt} be the optimal induced revenue curve. The additive revenue loss of ALG with respect to OPT is at most:*

$$n \cdot \max_{q \in [0,1]} \left(R^{opt}(q) - R^{alg}(q) \right).$$

Proof: First, to calculate the ex ante expected revenue of a bidder i with revenue curve R, ironing intervals $\mathcal{I}$ and reserve price r, we have by Myerson [27]:

$$Rev[R, \mathcal{I}, r] = -\mathbb{E}_{q \sim U[0,1]}[R(q) \cdot (y^{(\mathcal{I},r)})'(q)]. \tag{5}$$

Next, we apply the Switching Trick of Proposition 2.1. Let $\mathcal{I}_{alg}, \mathcal{I}_{opt}$ be the sets of ironing intervals and r_{alg}, r_{opt} be the reserve prices of ALG and OPT respectively. This yields total revenues:

$$\frac{Rev[F, I_{alg}, r_{alg}]}{n} = -\int_0^1 R^{alg}(q) y'(q) dq, \qquad \frac{Rev[F, I_{opt}, r_{opt}]}{n} = -\int_0^1 R^{opt}(q) y'(q) dq.$$

Note that the interim allocation function y in both cases is the same one; the only difference between $y^{(\mathcal{I}_{opt}, r_{opt})}$ and $y^{(\mathcal{I}_{alg}, r_{alg})}$ was the ironing intervals and reserve price, so after applying the switching trick, y is simply the truthful welfare-maximizing interim allocation rule (for single-item auctions this is the probability for bidder i to have the largest value $v_i = \max_j v_j$). This is the key point in our analysis, and it allows us to compare the expected revenue of both auctions directly:

$$\begin{aligned}
Rev[F, \mathcal{I}_{opt}, r_{opt}] - Rev[F, \mathcal{I}_{alg}, r_{alg}] &= n \left(-\int_0^1 R^{opt}(q) y'(q) dq + \int_0^1 R^{alg}(q) y'(q) dq \right) \\
&= n \int_0^1 \left(R^{opt}(q) - R^{alg}(q) \right) \left(-y'(q) \right) dq \\
&\leq n \cdot \max_{q \in [0,1]} \left(R^{opt}(q) - R^{alg}(q) \right) \cdot \int_0^1 -y'(q) dq \\
&= n \cdot \max_{q \in [0,1]} \left(R^{opt}(q) - R^{alg}(q) \right) \cdot (-y(1) + y(0)) \\
&= n \cdot \max_{q \in [0,1]} \left(R^{opt}(q) - R^{alg}(q) \right).
\end{aligned}$$

The inequality holds as $-y'$ is non-negative. The last equality holds as the probability of winning when your quantile is 0 is $y(0) = 1$, and the probability of winning when your quantile is 1 is $y(1) = 0$. Rearranging the terms yields the claim. ■

This is significant progress: the additive loss in revenue can be bounded in terms of the induced revenue curves of an algorithm ALG and the optimal algorithm, two objects that we have some hope of getting a handle on. Of course, we still need to show that the ironed revenue curve of Algorithm 2 is pointwise close to the ironed revenue curve induced by the optimal auction (Section 3.3).

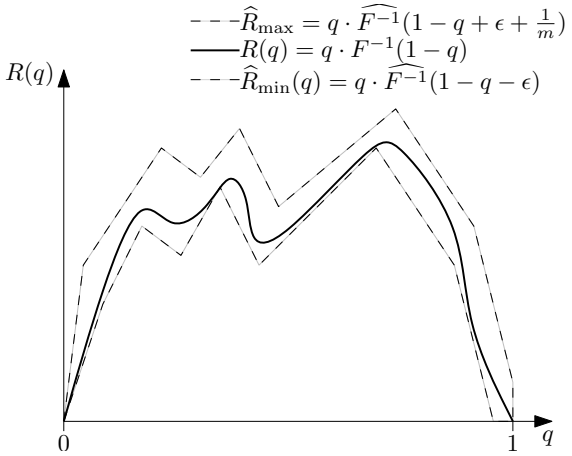

Figure 2: We can sandwich R between two estimated revenue curves $\widehat{R}_{\min}$ and $\widehat{R}_{\max}$.

3.3 Bounding the Error in the Revenue Curve

We implement the following steps to prove that the error in the learning algorithm's estimation of the revenue curve is small.

- (Lemma 3.3) We show that we can sandwich the actual revenue curve (without ironing or reserve price) R between two empirical revenue curves, $\widehat{R}_{\min}$ and $\widehat{R}_{\max}$ that are defined using the empirical quantile function.
- (Lemma 3.4 and Lemma 3.5) Let $\widehat{R}^{\star}_{\max}$ (resp. $\widehat{R}^{\star}_{\min}$) be the optimally induced revenue curve for $\widehat{R}_{\max}$ (resp. $\widehat{R}_{\min}$). The revenue curve induced by Algorithm 2, R^{alg}, is pointwise higher than the optimal induced revenue curve of the lower bound $\widehat{R}^{\star}_{\min}$, and the optimal induced revenue curve for the upper bound, $\widehat{R}^{\star}_{\max}$, is pointwise higher than R^{opt}.
- (Lemma 3.6) Finally, we show that $\widehat{R}^{\star}_{\max}(q) - \widehat{R}^{\star}_{\min}(q)$ is small for all q, and therefore the additive loss is small.

Lemma 3.3 *Let $\widehat{R}_{min}(q) = q \cdot \widehat{F^{-1}}(1-q-\epsilon)$ and $\widehat{R}_{max}(q) = q \cdot \widehat{F^{-1}}(1-q+\epsilon+\frac{1}{m})$. With probability at least $1-\delta$ for all $q \in [0,1]$:*

$$\widehat{R}_{min}(q) \quad \leq \quad R(q) \quad \leq \quad \widehat{R}_{max}(q).$$

Proof: See Figure 2 for graphical intuition. We will start with the first inequality: $\widehat{R}_{\min}(q) \leq R(q)$. By the DKW inequality with probability $1-\delta$ the following holds for all v:

$$-\epsilon \leq \widehat{F}(v) - F(v) \leq \epsilon \tag{6}$$

Define $q' = F(v)$. We can rewrite the first inequality of (6):

$$q' - \epsilon \leq \widehat{F}(v)$$
$$\widehat{F^{-1}}\left(q' - \epsilon\right) \leq \widehat{F^{-1}}\left(\widehat{F}(v)\right)$$

Since $\widehat{F^{-1}}$ is mononotonically non-decreasing. We now invoke (2):

$$\widehat{F^{-1}}\left(q' - \epsilon\right) \leq \widehat{F^{-1}}\left(\widehat{F}(v)\right) \leq v = F^{-1}(q').$$

Now let $q = 1 - q'$, and multiply both sides by q to obtain: $\widehat{R}_{\min}(q) \leq R(q)$.

The proof for the upper bound of R is analogous with the exception that we pick up another $\frac{1}{m}$ term to invoke (2). ■

So while the the algorithm does not know the exact revenue curve R, it can be upper bounded by $\widehat{R}_{\max}$ and lower bounded by $\widehat{R}_{\min}$. We'll use $\widehat{R}_{\min}$ to give a lower bound on the revenue curve R^{alg} induced by Algorithm 2, and $\widehat{R}_{\max}$ to give an upper bound on the revenue curve R^{opt} induced by Myerson's optimal auction. We start with the latter.

Lemma 3.4 *Let R^{opt} be the optimal induced revenue curve of R, and let $\widehat{R}^{\star}_{max}$ be the optimal induced revenue curve for $\widehat{R}_{max}$. Then with probability $1 - \delta$ for all $q \in [0, 1]$:*

$$R^{opt}(q) \quad \leq \quad \widehat{R}^{\star}_{max}(q).$$

Proof: By Lemma 3.3 we know that for every q, $\widehat{R}_{\max}(q) \geq R(q)$. First take the effect of optimally ironing both curves into account. Optimal ironing will lead to induced revenue curves $\mathcal{CH}(\widehat{R}_{\max})$ and $\mathcal{CH}(R)$ and since $\widehat{R}_{\max}$ is pointwise higher, so is its convex hull. What remains to be proven is that this property is maintained after the effect of the reserve price of the curve. The reserve quantile is the highest point on the curve, and the induced revenue curve will stay constant for all higher quantiles. Let q be the reserve quantile for R. If the reserve quantile for $\widehat{R}_{\max}$ happens to be the same, then the statement is proven. Moreover, if the optimal reserve quantile for $\widehat{R}_{\max}$ is not q, it can only be higher, in which case the statement is also proven. ■

So $\widehat{R}^{\star}_{\max}$ is pointwise higher than R^{opt}. Proving that $\widehat{R}^{\star}_{\min}$ is a lower bound for R^{alg} is slightly more involved since the ironing intervals and reserve price are given by Algorithm 2, and may not be optimal. Therefore, the induced revenue curve R^{alg} is in general not concave, and the reserve quantile may not be at the highest point of the curve. However, since R^{alg} is induced by the reserve price and ironing intervals that were chosen by looking at $\widehat{R}_{\min}$, there is enough structural information to prove the lower bound.

Lemma 3.5 *Let R^{alg} be the revenue curve induced by Algorithm 2 and let $\widehat{R}^{\star}_{min}$ be the optimal induced revenue curve for $\widehat{R}_{min}$. Then with probability $1 - \delta$ for all $q \in [0, 1]$:*

$$\widehat{R}^{\star}_{min}(q) \quad \leq \quad R^{alg}(q).$$

Proof: Again we first argue that after ironing the induced curve $\widehat{R}^{\star}_{\min}$ is completely below R^{alg}, and subsequently show that this is still true after setting the reserve price.

Algorithm 2 picks the ironing intervals based on where $\widehat{R}^{\star}_{\min}$ differs from $\mathcal{CH}(\widehat{R}_{\min})$. However, $\widehat{R}_{\min}$ and R^{alg} are defined in quantile space, whereas the ironing intervals are defined in value space. This means that R^{alg} is not necessarily ironed on the same intervals in quantile space as $\widehat{R}^{\star}_{\min}$. The line that goes through the origin and intersects $\widehat{R}_{\min}$ at q, will intersect R at the point q' where the $F^{-1}(q') = \widehat{F^{-1}}(1 - q - \epsilon)$, see Figure 3.[11] Any line through the origin intersects $\widehat{R}_{\min}$

[11] It's important to note that we do not need to know F^{-1}, to predict the effect of ironing the actual revenue curve.

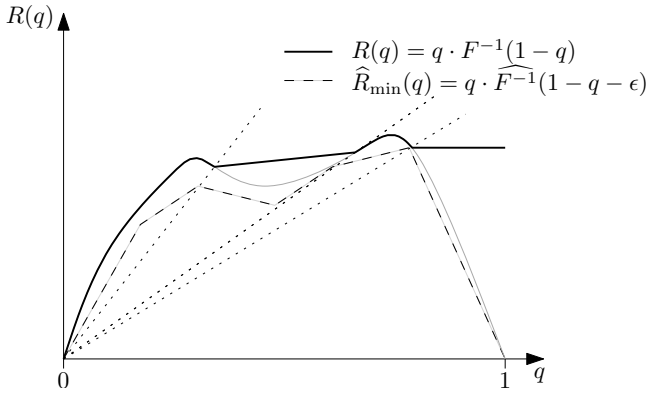

Figure 3: When we pick a value range to iron based on $\widehat{R}_{\min}$, its effect on the actual revenue curve can be seen. The quantiles of the start and end point of the ironing procedure are given by the line that intersects the start and end point on $\widehat{R}_{\min}$ and the origin.

before it intersects R, and since ironing is done according to these lines through the origin, this property is maintained after ironing. The reserve price is determined in the same way, and hence $R^{alg}(q) \geq \widehat{R}^{\star}_{\min}(q)$ for all $q \in [0,1]$ as claimed. ■

So we have our upper bound and lower bounds in terms of $\widehat{F^{-1}}$. Finally we show that the difference between the two is small.

Lemma 3.6 *For a distribution F, let R^{alg} be the revenue curve induced by Algorithm 2, R^{opt} be the optimal induced revenue curve. Using m samples from F, with probability $1-\delta$ and $\epsilon = \sqrt{\frac{\ln 2 \cdot \delta^{-1}}{2m}}$ for all q:*

$$R^{opt}(q) - R^{alg}(q) \leq \widehat{R}^{\star}_{max}(q) - \widehat{R}^{\star}_{min}(q) \leq \left(2\epsilon + \frac{1}{m}\right) H \leq 3\epsilon \cdot H.$$

Proof: Let $\mathcal{I}_{\max}$ and $r_{\max}$ be the set of ironing intervals and reserve price in quantile space of $\widehat{R}^{\star}_{\max}$. Comparing $\widehat{R}^{\star}_{\max}$ and $\widehat{R}^{\star}_{\min}$ directly is difficult since the ironing intervals and reserve price may be quite different between the two. Instead, we will take the set $\mathcal{I}_{\max}$ of ironing intervals $[a_i, b_i)$, and reserve quantile $r_{\max}$ from $\widehat{R}^{\star}_{\max}$ and use this to iron $\widehat{R}_{\min}(q)$ at intervals $[a_i - 2\epsilon - \frac{1}{m}, b_i - 2\epsilon - \frac{1}{m})$, and set a reserve price of $r_{\max} - 2\epsilon - \frac{1}{m}$.

If we compare $\widehat{R}^{\star}_{\max}$ against this induced curve $\widehat{R}^{(\mathcal{I}_{\max}, r_{\max})}_{\min}$, we have an upper bound for the difference with respect to $\widehat{R}^{\star}_{\min}$, since the latter is the optimal induced revenue curve and therefore pointwise higher. We can reason about ironing in quantile space without loss of generality since both $\widehat{R}_{\max}$ and $\widehat{R}_{\min}$ use the same function $\widehat{F^{-1}}$ (alternatively think of it as ironing in value space based on $\widehat{R}_{\max}$, the line through the origin and $\widehat{R}_{\max}(q)$ intersects $\widehat{R}_{\min}$ at $q - 2\epsilon - \frac{1}{m}$).

There are 3 cases to handle: 1) if q falls in an ironing interval $[a_i, b_i) \in \mathcal{I}_{\max}$, 2) if the quantile q is higher than the reserve quantile $q \geq r_{\max}$ of $\widehat{R}_{\max}$, and 3) when neither of those cases apply. We start with the case for a q where $\widehat{R}^{\star}_{\max}(q) = \widehat{R}_{\max}$, i.e. q does not fall in an ironing interval and is smaller than the reserve quantile:

$$\begin{aligned}
\widehat{R}^{\star}_{\max}(q) &= q \cdot \widehat{F^{-1}}\left(1 - q + \epsilon + \frac{1}{m}\right) \\
&\leq \left(q - 2\epsilon - \frac{1}{m}\right) \cdot \widehat{F^{-1}}\left(1 - q + \epsilon + \frac{1}{m}\right) + \left(2\epsilon + \frac{1}{m}\right) H \\
&= \widehat{R}_{\min}\left(q - 2\epsilon - \frac{1}{m}\right) + \left(2\epsilon + \frac{1}{m}\right) H \\
&= \widehat{R}^{\star}_{\min}\left(q - 2\epsilon - \frac{1}{m}\right) + \left(2\epsilon + \frac{1}{m}\right) H \\
&\leq \widehat{R}^{\star}_{\min}(q) + \left(2\epsilon + \frac{1}{m}\right) H
\end{aligned}$$

where the first inequality holds because $\widehat{F^{-1}}(q) \leq H$ for all q and the last inequality follows because $\widehat{R}^{\star}_{\min}$ is monotonically non-decreasing. Rearranging yields the claim for q not in an ironing interval or reserved.

The other cases follow similarly: consider the case where $q \geq r_{\max}$.

$$\begin{aligned}
\widehat{R}^{\star}_{\max}(q) &= \widehat{R}_{\max}\left(r_{\max}\right) \\
&= r_{\max} \cdot \widehat{F^{-1}}\left(1 - r_{\max} + \epsilon + \frac{1}{m}\right) \\
&\leq \left(r_{\max} - 2\epsilon - \frac{1}{m}\right) \cdot \widehat{F^{-1}}\left(1 - r_{\max} + \epsilon + \frac{1}{m}\right) + \left(2\epsilon + \frac{1}{m}\right) H \\
&= \widehat{R}_{\min}\left(r_{\max} - 2\epsilon - \frac{1}{m}\right) + \left(2\epsilon + \frac{1}{m}\right) H \\
&\leq \widehat{R}^{\star}_{\min}\left(r_{\max} - 2\epsilon - \frac{1}{m}\right) + \left(2\epsilon + \frac{1}{m}\right) H \\
&\leq \widehat{R}^{\star}_{\min}\left(q\right) + \left(2\epsilon + \frac{1}{m}\right) H
\end{aligned}$$

where the last inequality holds because $r_{\max} - 2\epsilon - \frac{1}{m} \leq r_{\max} \leq q$ and $\widehat{R}^{\star}_{\min}$ is non-decreasing.

Finally the case for when q falls in an ironing interval is analogous: $\widehat{R}^{\star}_{\max}(q)$ is a convex combination of the end points of the ironing interval: $\widehat{R}^{\star}_{\max}(q) = \frac{q-a_i}{b_i-a_i} \cdot \widehat{R}_{\max}(a_i) + \left(1 - \frac{q-a_i}{b_i-a_i}\right) \widehat{R}_{\max}(b_i)$ and hence

$$\widehat{R}^{\star}_{\min}(q) \geq \frac{q - a_i}{b_i - a_i} \cdot \widehat{R}_{\min}\left(a_i - 2\epsilon - \frac{1}{m}\right) + \left(1 - \frac{q - a_i}{b_i - a_i}\right) \widehat{R}_{\min}\left(b_i - 2\epsilon - \frac{1}{m}\right) + \left(2\epsilon + \frac{1}{m}\right) H.$$

And so $\widehat{R}^{\star}_{\max}(q) - \widehat{R}^{\star}_{\min}(q) \leq \left(2\epsilon + \frac{1}{m}\right) H$ in all cases. ■

Theorem 3.1 now follows by combining Lemmas 3.2 and 3.6. The additive loss the in expected revenue of Algorithm 2 is at most $3\epsilon \cdot n \cdot H$.

Proof of Theorem 3.1: By Lemma 3.2 we can express the total additive error of the expected revenue of an algorithm that yields ironing intervals $\mathcal{I}_{alg}$ and reserve price r_{alg} with respect to the

optimal auction as:

$$\frac{Rev[F, \mathcal{I}_{opt}, r_{opt}] - Rev[F, \mathcal{I}_{alg}, r_{alg}]}{n} \leq \max_{q \in [0,1]} \left(R^{opt}(q) - R^{alg}(q) \right).$$

By Lemma 3.6, Algorithm 2 yields

$$\max_{q \in [0,1]} \left(R^{opt}(q) - R^{alg}(q) \right) \leq \left(2\epsilon + \frac{1}{m} \right) H.$$

The theorem follows. ■

When valuations lie in $[1, H]$, so that the optimal expected revenue is bounded below, these results easily imply a sample complexity upper bound $O\left(\epsilon^{-2} \ln \epsilon^{-1} n^2 H^2\right)$ for learning (efficiently) a $(1-\epsilon)$-(multiplicative) approximate auction.

In Section 4 we show how to extend these results from single-item auctions to matroid environments and position environments.

4 Matroid and Position Environments

In Lemma 3.2 we showed that for environments with optimal auctions in $\mathcal{A}_{\mathcal{I},r}$, the additive loss of Algorithm 2 can be bounded in terms of the error in estimating the ironed revenue curve.

In this section we show that Myerson's optimal auction can be expressed in this way for more general single-parameter environments with i.i.d. bidders as well: namely in position auctions and auctions with matroid constraints. The theorem statement is as follows; the proof follows from Fact 4.2 and Fact 4.3 that show that the optimal auctions for those environments are in $\mathcal{A}$ (i.e., have the form $A^{(\mathcal{I},r)}$ for a suitable choice of ironed intervals $\mathcal{I}$ and reserve price r).

Theorem 4.1 *For position and matroid auctions with n i.i.d. bidders with values from unknown distribution F, m i.i.d. samples from F, with probability $1-\delta$, the additive loss in expected revenue of Algorithm 2 compared to the optimal expected revenue is at most $3 \cdot n \cdot H \cdot \sqrt{\frac{\ln 2\delta^{-1}}{2m}}$.*

Our results stand in contrast to previous work on picking the optimal reserve price based on samples, for which a $\Omega(\frac{\log n}{\log \log n})$ lower bound is known in these settings [11].

4.1 Position Auctions

A position auction [33] is one where the winners are given a position, and position i comes with a certain quantity x_i of the good. The canonical example is that of ad slot auctions for sponsored search, where the best slot has the highest click-through-rate, and subsequent slots have lower and lower click-through-rates. In an optimal auction, the bidder with the highest ironed virtual value gets the best slot, the second highest ironed virtual value the second slot, and so on.

Fact 4.2 *The optimal auction $A^{opt}_{(pos)}$ for position auctions can be expressed as an auction with ironing and reserve price in value space: $A^{opt}_{(pos)} \in \mathcal{A}$.*

Proof: In the optimal auction, the bidder with the highest ironed virtual value is awarded the first position (with allocation x_1), the bidder with the second highest ironed virtual value the second position with x_2, and so on. Since the ironed virtual value are monotonically non-decreasing in the value of a bidder, and identical in ironing intervals, this can equivalently be described by an auction in $\mathcal{A}$. ■

4.2 Matroid Environments

In a matroid environment, the feasible allocations are given by matroid $\mathcal{M} = (E, I)$, where E are the players and I are independent sets. The auction can simultaneously serve only sets S of players that form an independent set of the matroid $S \in I$. A special case of this is the rank k uniform matroid, which accepts all subsets of size at most k, i.e. it is a k-unit auction environment.

In matroid environments, the ex-post allocation function $x_i(\mathbf{b})$ and interim allocation function $y_i(q)$ are no longer the same for each player, e.g. imagine a player i who is not part of any independent set, then $y(q) = 0$ everywhere. So for matroid environments, the total additive loss is

$$\begin{aligned} Rev[F, \mathcal{I}_{opt}, r_{opt}] - Rev[F, \mathcal{I}_{alg}, r_{alg}] &= \sum_{i=1}^{n} \left(- \int_0^1 R^{opt}(q) y_i'(q) dq + \int_0^1 R^{alg}(q) y_i'(q) dq \right) \\ &= ... \\ &= n \cdot \max_{q \in [0,1]} \left(R^{opt}(q) - R^{alg}(q) \right). \end{aligned}$$

We need to show that the optimal allocation can still be expressed in terms of an auction $A \in \mathcal{A}$.

Fact 4.3 *The optimal auction $A^{opt}_{(mat)}$ for matroid auctions can be expressed as an auction with ironing and reserve price in value space: $A^{opt}_{(pos)} \in \mathcal{A}$.*

Proof of Fact 4.3: A property of matroids is that the following simple greedy algorithm yields the optimal solution:

GREEDY(E, I)
1 $S \leftarrow \emptyset$
2 **while** $\{i : i \notin S \wedge S \cup \{i\} \in I\} \neq \emptyset$
3 $\quad S \leftarrow \arg\max\{v_i : i \notin S \wedge S \cup \{i\} \in I\}$
4 **return** S

where v_i is the (ironed virtual) value associated with bidder i. Since the order of largest values is the same for both virtual values and bids (up to ties), the allocation of the optimal auction is identical to the auction that irons on $\mathcal{I}_{opt}$ and has reserve price r_{opt} (up to tie-braking); hence $A^{opt}_{(mat)} \in \mathcal{A}$. ■

5 No-Regret Algorithm

So far we assumed access to batch of samples before having to choose an auction. In this section we show that even if we start out without any samples, but run the auction is a repeated setting, using past bids as samples leads to a no-regret algorithm. The goal here is to achieve additive error

Algorithm 3 A no-regret algorithm for optimal auctions.

NO-REGRET-AUCTION(δ, T)

```
1  ▷ Round 0:
2  Collect a set of bids b, run VCG(b)
3  X ← b
4  for round t = 1...T
5        Collect a set of bids b
6        EMPIRICALMYERSON(X, δ/T, b)
7        X ← X ∪ b
```

$o(T) \cdot O(\text{poly}(n, H, \delta))$ — the error can be polynomial in all parameters except the time horizon T, for which it should be strictly sublinear. We show that for Algorithm 3 the total loss grows as $\widetilde{O}(\sqrt{T}\sqrt{n}\sqrt{\log(\delta^{-1})}H)$ and hence results in a no-regret algorithm.

Invoking Theorem 3.1 with confidence parameter δ/T and taking a union bound, we have the following.

Fact 5.1 *With probability $1-\delta$, for all rounds simultaneously, each round $t \in [1, T]$ of Algorithm 3 has additive loss at most $3\sqrt{\frac{\ln(2T\delta^{-1})}{2nt}} \cdot n \cdot H$.*

This leads to the following no-regret bound.

Theorem 5.2 *With probability $1-\delta$, the total additive loss of Algorithm 3 is $O(\sqrt{n}\sqrt{T \log T}\sqrt{\log \delta^{-1}} \cdot H)$, which is $\widetilde{O}(T^{1/2})$ with respect to T.*

Proof: By Fact 5.1 with probability $1-\delta$ for all rounds simultaneously, the additive loss for round t is bounded by $3\sqrt{\frac{\ln(2T\delta^{-1})}{2nt}} \cdot n \cdot H$. The loss of day 0 is at most $H \cdot n$. The total loss can then be bounded by:

$$(n \cdot H) \cdot \left(1 + 3\sum_{t=1}^{T} \sqrt{\frac{\ln(2T\delta^{-1})}{2nt}}\right)$$

We can rewrite the sum:

$$\begin{aligned}\sum_{t=1}^{T} \sqrt{\frac{\ln(2T/\delta)}{2nt}} &= \sqrt{\frac{\ln(2T/\delta)}{2n}} \cdot \sum_{t=1}^{T} \sqrt{\frac{1}{t}} \\ &\leq \sqrt{\frac{\ln(2T/\delta)}{2n}} \cdot 2\sqrt{T} \\ &= \sqrt{\frac{2T\ln(2T/\delta)}{n}}\end{aligned}$$

Hence the total loss is

$$(n \cdot H) \cdot \left(1 + 3\sqrt{\frac{2T\ln(4T/\delta)}{n}}\right) = O(\sqrt{n}\sqrt{T\log T}\sqrt{\log \delta^{-1}} \cdot H);$$

the dependence on T is $O(\sqrt{T \log T}) = \widetilde{O}(\sqrt{T})$. ■

The bound of $O(\sqrt{T \log T})$ is almost tight, as there is a lower bound of $\Omega(\sqrt{T})$ given by [6]. Also note that if we do not know T a priori, we can use a standard doubling argument to obtain the same asymptotic guarantee.

References

[1] Pablo Daniel Azar, Constantinos Daskalakis, Silvio Micali, and S. Matthew Weinberg. Optimal and efficient parametric auctions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 596–604, 2013.

[2] Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting prices with unknown distributions. In *Innovations in Computer Science (ICS)*. Tsinghua University Press, January 2011.

[3] Maria-Florina Balcan, Avrim Blum, Jason D. Hartline, and Yishay Mansour. Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci.*, 74(8):1245–1270, 2008.

[4] S. Baliga and R. Vohra. Market research and market design. *Advances in Theoretical Economics*, 3, 2003. Article 5.

[5] N. Cesa-Bianchi, C. Gentile, and Y. Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, 2015.

[6] Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Regret minimization for reserve prices in second-price auctions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 1190–1204. SIAM, 2013.

[7] Shuchi Chawla, Jason D. Hartline, and Denis Nekipelov. Mechanism design for data science. *CoRR*, abs/1404.5971, 2014.

[8] Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu. Mechanism design with approximate valuations. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 34–38, 2012.

[9] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 243–252. ACM, 2014.

[10] Jacques Crémer and Richard P. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2):345–361, 1985.

[11] Nikhil R Devanur, Jason D Hartline, and Qiqi Yan. Envy freedom and prior-free mechanism design. *Journal of Economic Theory*, 2014.

[12] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, pages 129–138, 2010.

[13] S. Dughmi, L. Han, and N. Nisan. Sampling and representation complexity of revenue maximization. In *Workshop on Internet and Network Economics (WINE)*, pages 277–291, 2014.

[14] Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, pages 642–669, 1956.

[15] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[16] Edith Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 736–745. SIAM, 2007.

[17] Hu Fu, Nima Haghpanah, Jason D. Hartline, and Robert Kleinberg. Optimal auctions for correlated buyers with sampling. In *EC*, pages 23–36, 2014.

[18] Andrew V Goldberg, Jason D Hartline, Anna R Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

[19] Bach Q Ha and Jason D Hartline. Mechanism design via consensus estimates, cross checking, and profit extraction. *ACM Transactions on Economics and Computation*, 1(2):8, 2013.

[20] Jason Hartline. Mechanism design and approximation. Book draft (September 2014), September 2014.

[21] Jason D Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 75–84. ACM, 2008.

[22] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. *Proceedings of EC'15*, 2015.

[23] Robert D. Kleinberg and Frank Thomson Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS*, pages 594–605, New York, New York, USA., 2003. IEEE Computer Society.

[24] Pascal Massart. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability*, pages 1269–1283, 1990.

[25] Andres Munoz Medina and Mehryar Mohri. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of The 31st Intl. Conf. on Machine Learning*, pages 262–270, 2014.

[26] J. Morgenstern and T. Roughgarden. The psuedo-dimension of near-optimal auctions. To appear in NIPS '16, 2015.

[27] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[28] Z. Neeman. The effectiveness of English auctions. *Games and Economic Behavior*, 43(2):214–238, 2003.

[29] Noam Nisan. Introduction to mechanism design (for computer scientists). In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[30] I. Segal. Optimal pricing mechanisms with unknown demand. *American Economic Review*, 93(3):509–529, 2003.

[31] Balasubramanian Sivan and Vasilis Syrgkanis. Vickrey auctions for irregular distributions. In *Web and Internet Economics*, pages 422–435. Springer, 2013.

[32] Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

[33] Hal R Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.

A Reduced Information Model

This appendix considers a reduced information model for i.i.d. bidders where the auctioneer can only set a reserve price and then observes the selling price; see also [6]. The goal is to model an observer who can see the outcome of past auctions, and perhaps submit a "shill bid" to set a reserve, but cannot directly observe the bids.

We assume we observe m samples from the second highest order statistic (so out of n i.i.d. bids, we see the second highest bid). We show that there are distributions such that, in order to get the performance close to that of running Myerson in the reduced information model, you first need to see at least an exponential number of samples. This is far more than what would be needed with regular valuation distributions, where only the monopoly price is relevant.

A.1 Lower Bound

Theorem A.1 *For any $\epsilon > 0$, to obtain $\frac{1}{4} - \epsilon$ additive loss compared to Myerson's auction, with constant probability you need $\Omega(2^n/n)$ samples from $F_{(2)}$.*

Before proving the statement, let's think about what this means: it means that if all we observe are samples from $F_{(2)}$, for large n, we cannot hope to get a vanishing regret in a polynomial (in n) number of steps. Moreover, on this distribution, Myerson obtains at most 2 expected revenue, so we lose $\frac{1}{8}$ of the profit (see corollary after the proof).

Proof of Theorem A.1: We'll give 2 distributions, D_1 and D_2, such that we need $\Omega(2^n)$ samples from $F_{(2)}$ to differentiate between the two. Moreover, we'll show that every auction that approximates the optimal auction for either D_1 or D_2, incurs an additive loss of at least $\frac{1}{2}(1 - \frac{1}{\sqrt{2}})$ on one of the two.

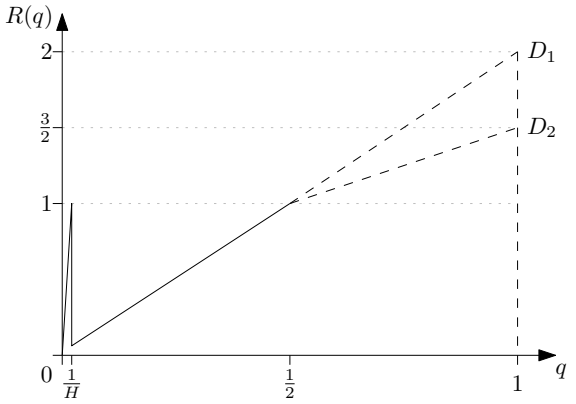

(a) Two distributions, D_1 and D_2 that are hard to distinguish using samples from $F_{(2)}$. The distributions agree on $q \in [0, 1/2]$ and disagree elsewhere.

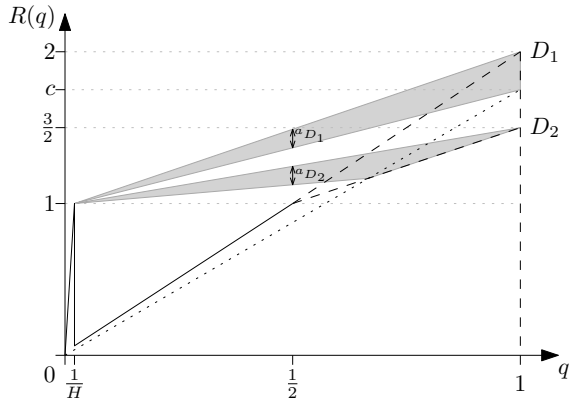

(b) Consider ironing on $[c, H)$. The top shaded area indicates the loss with respect to D_1 and the bottom shaded area represents the loss with respect to D_2.

Figure 4: Lower bound example for optimal auctions using samples from $F_{(2)}$.

We define D_1 and D_2 by their quantile function:

$$D_1^{-1}(q) = \begin{cases} H & \text{for } q \leq \frac{1}{H} \\ 2 & \text{otherwise} \end{cases}$$

$$D_2^{-1}(q) = \begin{cases} H & \text{for } q \leq \frac{1}{H} \\ 2 & \text{for } \frac{1}{H} < q \leq \frac{1}{2} \\ 1 + \frac{1}{2q} & \text{otherwise.} \end{cases}$$

Here $H \gg n$ is a sufficiently large constant. The two distributions agree on $q \in [0, \frac{1}{2}]$. For $q \in (\frac{1}{2}, 1]$ they differ, and the effect on the revenue curve can be seen in Figure 4a.

To complete the proof we need to show 2 things: 1) that differentiating between D_1 and D_2 based on samples from $F_{(2)}$ requires $\Omega(2^n/2)$ samples, and 2) that no auction can simultaneously approximate the optimal solution for both.

The first aspect of this is straightforward. Whenever we observe a sample from $F_{(2)}$ from the top half quantile: $q \in [0, \frac{1}{2}]$, we get no information, since this is the same for both distributions. So a necessary condition for differentiating between D_1 and D_2 is if we observe a sample from $q \in (\frac{1}{2}, 1]$. For this to happen, we need that out of n draws, $n-1$ draws are in the bottom quantile. This happens with probability $\frac{n}{2^{n-1}}$, so after $\frac{2^n}{n}$ samples, with probability approximately $1 - \frac{1}{e}$ we haven't seen a sample that will differentiate the two distributions.

Now we will show that until the moment when you can differentiate between D_1 and D_2, there is no way to run an auction that performs well for both. Note that the optimal auction for D_1 is to iron $[2, H)$ and the optimal auction for D_2 is to iron $[\frac{3}{2}, H)$. Neither has a reserve price. The set of all auctions that could potentially work well for either, is the set of auctions that irons $[c, H)$ (see Figure 4b).

We'll only count the loss we incur in the range $[0, \frac{1}{2}]$, take $H \to \infty$ and we see that if a is the height of the shaded region (of D_1 resp. D_2) at $q = \frac{1}{2}$, then its loss with respect to the optimal auction is:

$$
\begin{aligned}
n \cdot \int_0^{\frac{1}{2}} 2aq \cdot (n-1)(1-q)^{n-2} dq &= \left[2anq(1-q)^{n-1}\right]_0^{1/2} - 2an \int_0^{\frac{1}{2}} (1-q)^{n-1} dq \\
&= 2an \left(\frac{1}{2}\right)^{n-1} - 2an \left[\frac{1}{n}(1-q)^n\right]_0^{1/2} \\
&= 2an \left(\frac{1}{2}\right)^{n-1} - 2an \frac{1}{n} \left(\frac{1}{2}\right)^n + 2an \frac{1}{n} \\
&= 2an \left(\frac{1}{2}\right)^{n-1} \left(1 - \frac{1}{2n}\right) + 2a \\
&\geq 2a
\end{aligned}
$$

This means, that if we can show that for all choices of $c \in [3/2, 1]$ at least one of D_1, D_2 has a large a, we are done. The closer c is to 2, the smaller a_{D_2} is, the closer c is to $3/2$, the smaller a_{D_1}, so we'll balance the two and show that neither is small enough.

Finding a_{D_1} in terms of c is easy enough: $a_{D_1} = \frac{2-c}{2}$. a_{D_2} is slightly harder: we first need to find the intersection of $q \cdot c$ and $R_{D_2}(q)$:

$$
\begin{aligned}
\frac{1}{2} + q &= qc \\
(c-1)q &= \frac{1}{2} \\
q &= \frac{1}{2(c-1)}
\end{aligned}
$$

So the lower part of the shaded area of D_2 is a line that passes through the points $(0, 1)$ and $(\frac{1}{2(c-1)}, \frac{c}{2(c-1)})$. The line segment that connects the two is given by the equation $1 + q(2-c)$ hence at $q = \frac{1}{2}$ it is $2 - c/2$, therefore $a_{D_2} = \frac{5}{4} - 2 + \frac{c}{2} = \frac{c}{2} - \frac{3}{4}$.

Setting $a_{D_1} = a_{D_2}$ yields $c = \frac{7}{4}$ with $a_{D_1} = a_{D_2} = \frac{1}{8}$. Therefore, for any auction that we decide to run will have additive loss of $2a = \frac{1}{4}$ on one of the two distributions, and we need $\Omega(2^n/n)$ samples to decide which distribution we are dealing with. ■

Corollary A.2 *For any $\epsilon > 0$, to obtain a $(\frac{7}{8} + \epsilon)$-approximation (multiplicative) to Myerson's auction, with constant probability you need $\Omega(2^n/n)$ samples.*

Proof: We can upper bound the revenue curve by the constant 2, to show that the optimal auction cannot have expected revenue more than 2. Since we have additive loss of $1/4$, the multiplicative approximation ratio is at most $\frac{7}{8}$. ■